

# ADAPTIVE HAMILTONIAN ESTIMATION USING BAYESIAN EXPERIMENTAL DESIGN

Christopher Ferrie<sup>\*,†</sup>, Christopher E. Granade<sup>\*,\*\*</sup> and D. G. Cory<sup>‡,\*,§</sup>

<sup>\*</sup>*Institute for Quantum Computing, University of Waterloo, Waterloo, Ontario, Canada*

<sup>†</sup>*Department of Applied Mathematics, University of Waterloo, Ontario, Canada*

<sup>\*\*</sup>*Department of Physics, University of Waterloo, Ontario, Canada*

<sup>‡</sup>*Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada*

<sup>§</sup>*Department of Chemistry, University of Waterloo, Ontario, Canada*

**Abstract.** Using Bayesian experimental design techniques, we have shown that for a single two-level quantum mechanical system under strong (projective) measurement, the dynamical parameters of a model Hamiltonian can be estimated with exponentially improved accuracy over offline estimation strategies. To achieve this, we derive an adaptive protocol which finds the optimal experiments based on previous observations. We show that the risk associated with this algorithm is close to the global optimum, given a uniform prior. Additionally, we show that sampling at the Nyquist rate is not optimal.

**Keywords:** Experimental design, adaptive, parameter estimation, quantum, tomography

**PACS:** 07.05.Fb, 07.05.Kf, 03.65.Aa, 03.67.-a, 03.65.Ta

## INTRODUCTION

Quantum mechanics gives the most accurate description of many physical systems of interest. In turn, the most accurate characterization of a quantum device is given by its quantum mechanical model. Thus, efficient methods for the honest estimation of the distribution of parameters in a quantum mechanical model are of utmost importance, not only for building robust quantum technologies, but to reach new regimes of physics.

Bayesian experimental design (see, e.g. [1]) is a methodology to ascertain the utility of a proposed experiment. Bayesian experimental design has been successfully applied to problems in experimental physics, such as in the recent examples of [2] and [3]. In classical theories of physics and statistics, the measurement simply reveals the state of the system at that instant. By contrast, quantum theory presents with the following physical (and conceptual) barrier: no single measurement can reveal the state. Rather, each potential kind of experiment admits a probability distribution from which we draw our data. Thus, the methodology of experimental design seems tailor-made for quantum theory.

The structure of the paper is as follows. We begin by reviewing the general outline of Bayesian experimental design. We then apply the technique to devise an algorithm for the estimation of quantum Hamiltonian parameters. We show that in a particular case, this strategy is nearly globally optimal and demonstrate its improvement over standard algorithms numerically. Finally we conclude with a discussion on the applicability of this technique to real experiments on more complex quantum systems.

## BAYESIAN EXPERIMENTAL DESIGN

We assume some initial experiment  $E$  has been performed and data  $D$  has been obtained. The goal is to determine  $\Pr(\Theta|D, E)$ , the probability distribution of the model parameters  $\Theta$  given the experimental data. To achieve this we use Bayes' rule

$$\Pr(\Theta|D, E) = \frac{\Pr(D|\Theta, E) \Pr(\Theta|E)}{\Pr(D|E)},$$

where  $\Pr(D|\Theta, E)$  is the *likelihood function*, which is determined through the process of modeling the experiment, and  $\Pr(\Theta|E)$  is the *prior*, which encodes any *a priori* knowledge of the model parameters. The final term  $\Pr(D|E)$  can simply be thought as a normalization factor.

At this stage we can stop or obtain further data. Experimental design is well suited to quantum theory since an arbitrary fixed measurement procedure does not give maximal knowledge as is often assumed in the statistical modeling of classical system. We conceive, then, of possible future data  $D_1$  obtained from a, possibly different, experiment  $E_1$ . The probability of obtaining this data can be computed from the distributions at hand via marginalizing over model parameters

$$\Pr(D_1|E_1, D, E) = \int \Pr(D_1|\Theta, E_1) \Pr(\Theta|D, E) d\Theta.$$

We can use this distribution to calculate the expected *utility* of an experiment

$$U(E_1) = \sum_{D_1} \Pr(D_1|E_1, D, E) U(D_1, E_1),$$

where  $U(D_1, E_1)$  is the utility we would derive if experiment  $E_1$  gave result  $D_1$ . This could in principle be any function tailored to the specific problem. However, for scientific inference, a generally well motivated measure of utility is *information gain* [4]. In information theory, information is measured by the entropy

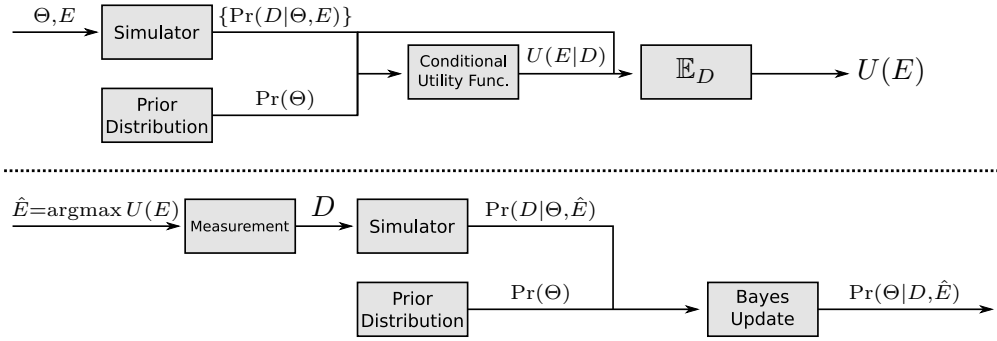
$$U(D_1, E_1) = \int \Pr(\Theta|D_1, E_1, D, E) \log \Pr(\Theta|D_1, E_1, D, E) d\Theta.$$

Thus, we search for the experiment which maximizes the expected information in the final distribution. That is, an optimal experiment  $\hat{E}$  is one which satisfies

$$U(\hat{E}) = \max_{E_1} \left\{ \sum_{D_1} \Pr(D_1|E_1, D, E) \times \int \Pr(\Theta|D_1, E_1, D, E) \log \Pr(\Theta|D_1, E_1, D, E) d\Theta \right\}.$$

## APPLICATION TO SIMPLE EXAMPLE

As an example of how to apply the Bayesian experimental design formalism to problems in quantum information, we consider a simple situation with a single qubit. In particular,



**FIGURE 1.** Overview of a step in the online adaptive algorithm for finding locally optimal experiments. Top: Method for calculating the utility function  $U(E)$ , given a simulator and a prior distribution  $\Pr(\Theta)$  over model parameters  $\Theta$ . Bottom: Method for updating prior distribution with results  $D$  from chosen actual experiment.

we suppose that the qubit evolves under an internal Hamiltonian

$$H = \frac{\omega}{2} \sigma_z.$$

Here  $\omega$  is an unknown parameter whose value we want to estimate. An experiment consists of preparing a single known input state  $\psi_{\text{in}} = |+\rangle$ , the  $+1$  eigenstate of  $\sigma_x$ , evolving under the Hamiltonian  $H$  for a controllable time  $t$  and performing a measurement in the  $\sigma_x$  basis. This is the simplest problem where adaptive Hamiltonian estimation can be used and is the problem studied in reference [5].

In the language of Bayesian inference, the data  $D \in \{0, 1\}$  is the outcome of the measurement. An experiment  $E$  consists of a specification of time  $t$  that the Hamiltonian is on, while the model parameter  $\Theta$  is simply  $\omega$ . The likelihood function is given by the Born rule

$$\Pr(D=0|\Theta, E) = |\langle + | e^{i\frac{\omega}{2}\sigma_z t} | + \rangle|^2 = \cos^2\left(\frac{\omega}{2}t\right).$$

Experimental design is a decision theoretic problem based on the utility function

$$U(t) = \sum_D \Pr(D|t) \int \Pr(\omega|D, t) \log \Pr(\omega|D, t) d\omega.$$

The optimal design is any value of  $t$  which maximizes this quantity.

We proceed by performing the optimal experiment and obtaining data  $D_1$ . Using Bayesian inference we update our prior  $\Pr(\omega)$  via Bayes' rule:

$$\Pr(\omega|D_1) = \frac{\Pr(D_1|\omega) \Pr(\omega)}{\Pr(D_1)}.$$

If we are not satisfied, we can repeat the process where this distribution becomes the prior for the new experimental design step. This algorithm is depicted in figure 1.

## Estimators, squared error loss and a greedy alternative to information gain

The preceding problem had a single unknown variable. If we desire an estimate  $\hat{\Theta}$  of the true value  $\Theta$ , the most often used figure of merit is the *squared error loss*:

$$L(\Theta, \hat{\Theta}) = |\Theta - \hat{\Theta}|^2.$$

The *risk* of an estimator  $\hat{\Theta} : \{D, D_1, D_2, \dots, D_N\} \mapsto \mathbb{R}$  is its expected performance with respect to the loss function:

$$R(\Theta, \hat{\Theta}) = \sum_{\{D, D_1, D_2, \dots, D_N\}} \Pr(\{D, D_1, D_2, \dots, D_N\} | \Theta) L(\Theta, \hat{\Theta}).$$

For squared error loss, the risk is also called the *mean squared error*. The average of this quantity with respect to some prior  $\Pr(\Theta) =: \pi(\Theta)$  is the *Bayes risk* of  $\pi$ ,

$$r(\pi, \hat{\Theta}) = \int R(\Theta, \hat{\Theta}) \pi(\Theta) d\Theta,$$

and the estimator which minimizes this quantity is called a *Bayes estimator*. In this case the Bayes estimator is the mean of the posterior distribution<sup>1</sup>. Let us assume then that the estimators we choose are Bayes. Let us also choose a uniform prior for  $\Theta$ . Then, the final figure of merit is the average mean squared error (AMSE):

$$r = \int R(\Theta, \hat{\Theta}) d\Theta.$$

We would like a strategy which minimizes this quantity. Non-adaptive Fourier and Bayesian strategies were investigated and compared to an adaptive strategy in reference [5]. Their adaptive strategy fits into the Bayesian experimental design framework when the utility is measured by the *variance* of the posterior distribution:

$$V(D_1, E_1) = - \int \Pr(\Theta | D_1, E_1, D, E) (\Theta^2 - \mu(D_1, E_1))^2 d\Theta,$$

where

$$\mu(D_1, E_1) = \int \Pr(\Theta | D_1, E_1, D, E) \Theta d\Theta$$

is the mean of the posterior. Recall that the mean is a Bayes estimator of AMSE, so  $\mu = \hat{\Theta}$ . For a single measurement this utility function satisfies  $V = -r$ . That is, maximizing the utility *locally* at each step of the algorithm is equivalent to minimizing the AMSE at each step. Hence, when using the negative variance as our utility function, the adaptive strategy summarized in Figure 1 is an example not only of a local optimization, but also

---

<sup>1</sup> Note that in any case where the loss function is *strictly proper*, i.e. is equal to zero if and only if the estimate is equal to the true state, the Bayes estimator is the posterior mean [6].

a *greedy algorithm* with respect to the AMSE risk. In the future, we shall refer to this choice of utility function together with the local optimization algorithm as the greedy algorithm for this problem.

We can write the risk of this strategy recursively as follows. Suppose at the  $N$ 'th, and final, measurement we have the updated distribution  $\pi_{N-1}$ . Then, the risk of the local strategy is

$$l_N(\pi_{N-1}, \Theta) = \sum_{D_N} \Pr(D_N | \Theta, \hat{E}_N) L(\Theta, \mu(D_N, \hat{E}_N)),$$

where  $\hat{E}_N$  is the locally optimal design satisfying

$$\hat{E}_N = \operatorname{argmin}_{E_N} \int \sum_{D_N} \Pr(D_N | \Theta, E_N) L(\Theta, \mu(D_N, E_N)) \pi_{N-1}(\Theta) d\Theta.$$

The expected risk at any other stage is

$$l_n(\pi_{n-1}, \Theta) = \sum_{D_n} \Pr(D_n | \hat{E}_n) l_{n+1} \left( \frac{\Pr(D_n | \Theta, \hat{E}_n) \pi_{n-1}(\Theta)}{\int \Pr(D_n | \Theta, \hat{E}_n) \pi_{n-1}(\Theta) d\Theta} \right),$$

where  $\hat{E}_n$  is, again, the locally optimal design satisfying

$$\hat{E}_n = \operatorname{argmin}_{E_n} \int \sum_{D_n} \Pr(D_n | \Theta, E_n) L(\Theta, \mu(D_n, E_n)) \pi_{n-1}(\Theta) d\Theta.$$

Then, the Bayes risk of the greedy strategy is

$$\int l_1(\pi_0, \Theta) \pi_0(\Theta) d\Theta.$$

Again, it is clear that the greedy algorithm is globally optimal on the final decision, as there is no further hypothetical data to consider. That is, the optimal solution at the  $N$ 'th measurement is

$$g_N(\pi_{N-1}, \Theta) = \sum_{D_N} \Pr(D_N | \Theta, \hat{E}_N) L(\Theta, \mu(D_N, \hat{E}_N)),$$

where  $\hat{E}_N$  is the locally optimal design satisfying

$$\hat{E}_N = \operatorname{argmin}_{E_N} \int \sum_{D_N} \Pr(D_N | \Theta, E_N) L(\Theta, \mu(D_N, E_N)) \pi_{N-1}(\Theta) d\Theta.$$

However, the globally optimal risk at any other stage

$$g_n(\pi_{n-1}, \Theta) = \sum_{D_n} \Pr(D_n | \tilde{E}_n) g_{n+1} \left( \frac{\Pr(D_n | \Theta, \tilde{E}_n) \pi_{n-1}(\Theta)}{\int \Pr(D_n | \Theta, \tilde{E}_n) \pi_{n-1}(\Theta) d\Theta} \right),$$

where now  $\tilde{E}_n$  is the globally optimal design satisfying

$$\tilde{E}_n = \operatorname{argmin}_{E_n} \int \sum_{D_n} \Pr(D_n | \Theta, E_n) g_{n+1} \left( \frac{\Pr(D_n | \Theta, E_n) \pi_{n-1}(\Theta)}{\int \Pr(D_n | \Theta, E_n) \pi_{n-1}(\Theta) d\Theta} \right) \pi_{n-1}(\Theta) d\Theta.$$

Then, the Bayes risk of the greedy strategy is

$$\int g_1(\pi_0, \Theta) \pi_0(\Theta) d\Theta.$$

In general,  $l_1(\pi_0, \Theta) \neq g_1(\pi_0, \Theta)$ . Nor is it the case that

$$\int l_1(\pi_0, \Theta) \pi_0(\Theta) d\Theta = \int g_1(\pi_0, \Theta) \pi_0(\Theta) d\Theta$$

for an arbitrary prior. However, for the special case of the uniform prior, we have found numerically that the Bayes risk of the greedy strategy and the Bayes risk of the global strategy are similar enough that the greedy strategy is useful.

## Performance comparisons

In reference [5], it was shown via simulation that the posterior variance of the greedy strategy is best fit by an exponentially decreasing function of  $N$ , the total number of measurements. In contrast, all off-line strategies decrease at best as a linear function of  $N$ .

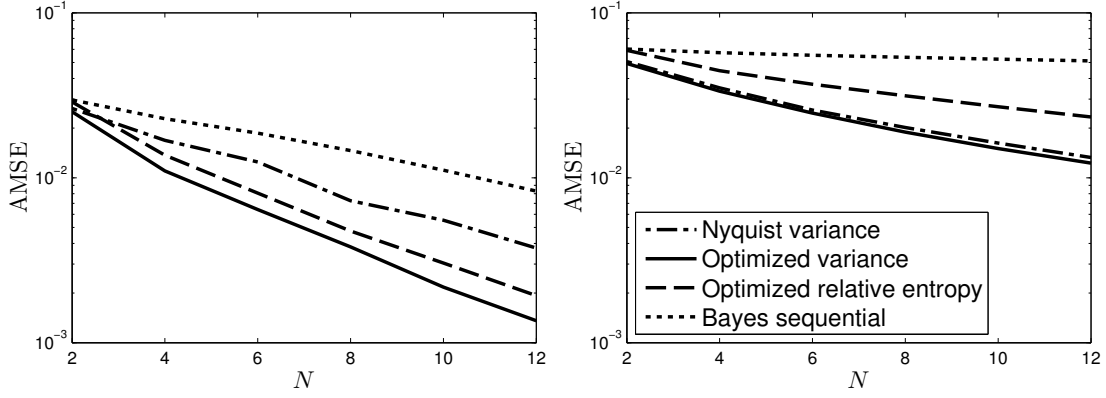
In Figure 2, we show that the local information gain optimizing algorithm also enjoys an exponential improvement in accuracy over naive off-line methods. Moreover, we show Nyquist rate sampling is unnecessary and, indeed, sub-optimal. All results stated are obtained using a uniform prior on  $[0, 1]$  and are computed numerically by exploring every branch of the the decision tree, in contrast to simulation.

In order to be “fair” to the off-line methods, we restricted the adaptive methods to explore the same experimental design specifications. That is, for this particular problem, the adaptive algorithm was allowed to select measurement times from  $[0, N_{\max}\pi]$ , where  $N_{\max}$  is the total number of measurements. In principle, these methods could only do better with a larger design specification.

## DISCUSSION

Summarizing, we have shown for the problem of estimating the parameter in a simple Hamiltonian model of qubit dynamics an adaptive measurement strategy can exponentially improve the accuracy over offline estimation strategies. Moreover, we have shown that sampling at the Nyquist rate is not optimal in the case of strong measurement. We have derived a recursive solution to the risks for both the local and global optimal strategies. Using this solution, we numerically found that the local strategy is nearly optimal in the special case of a uniform prior. That the greedy algorithm is nearly optimal in a case relevant to experiment demonstrates that an adaptive Bayesian method may be computationally feasible, in that an implementation need not consider all possible future data when choosing each experiment.

Together, these results demonstrate the usefulness of an adaptive Bayesian algorithm for parameter estimation in quantum mechanical systems, especially in comparison with



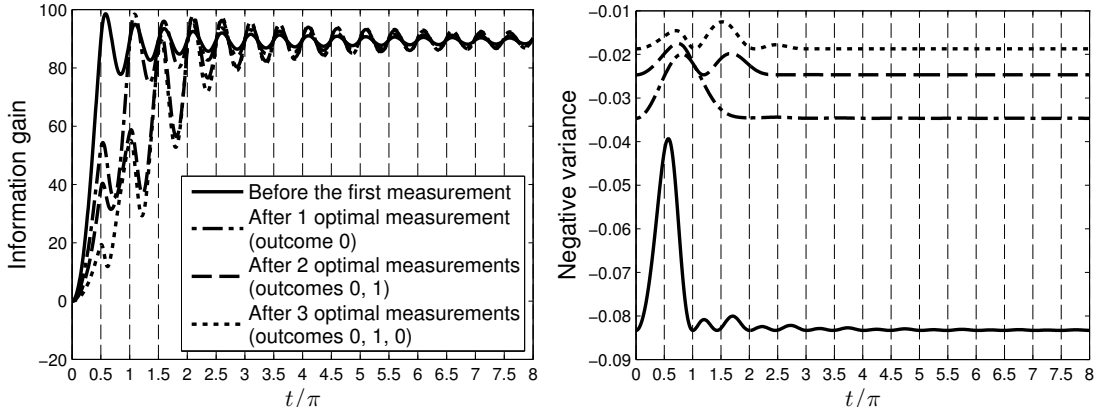
**FIGURE 2.** Performance of the estimation strategies. The Bayesian sequential and the strategy labeled “Nyquist” sample at the Nyquist rate. The “optimized” strategies find the global maximum utility (using Matlab’s “fmincon” starting with the optimal Nyquist time). In each case,  $N_{\max} = 12$  measurements are considered. Left: the ideal model discussed in the text. Right: a more realistic model with 25% noise and an addition relaxation process (known as  $T_2$ ) which exponentially decays the signal (to half its value at  $t = 10\pi$ ).

other algorithms in common use. In the presence of noise, this improvement becomes still more stark, as demonstrated by the results shown in Figure 2.

Why is it the case that the Nyquist times are not optimal? First, why should we expect them to be optimal? The Nyquist theorem states that a signal which contains no frequencies higher than  $\omega_{\max}$  is completely and unambiguously characterized by a discrete set of samples taken at a rate greater than or equal to  $\pi/2\omega_{\max}$ . However, the classical notion of sampling fails for the strong-measurement case that we consider here. What we have is a periodic *probability* distribution which can be sampled, not a periodic function whose *values* can be ascertained. That is there is no *signal*, in the classical sense of the word, which can be reconstructed. The failure of the Nyquist rate sampling is exemplified in Figure 3.

In this paper, we have chosen to measure success via the squared error loss. Although this is a standard metric, note that it is not practically useful in the context of estimating the parameters of a quantum mechanical system. We motivate this claim as follows. A typical application of our algorithm is to inform control theory algorithms, which can achieve significantly higher fidelities if given a distribution over Hamiltonians rather than a single best estimate. Indeed, in the case of nuclear magnetic resonance, the physical ensemble of qubits produces a real distribution of Hamiltonians to which control theory algorithm must be robust against [7, 8]. Any single estimate of the Hamiltonian parameters will thus artificially exclude dynamics which will appear as decoherence in the resultant pulses. Thus, we must measure the success of our algorithm via a loss function of the true *distribution* and estimated posterior. Noting that relative entropy is broadly considered the correct loss function for probability estimators, our algorithm, which maximizes expected information gain, becomes the optimal solution.

We expect that in more complicated systems, the Bayesian adaptive method will remain useful, especially in applications such as optimal control theory, where having a distribution over Hamiltonians is significantly more useful than a single best estimate.



**FIGURE 3.** The information gain (left) and variance (right) utilities for the prior followed by three simulated measurements. The vertical grid lines indicate the Nyquist times. Note that the times at which the utilities are maximized do not necessarily increase with the number of measurements.

## ACKNOWLEDGMENTS

CF thanks Josh Combes for helpful discussions. This work was financially supported by NSERC and CERC.

## REFERENCES

1. T. J. Lored, *AIP Conference Proceedings* **707**, 330–346 (2004), URL <http://dx.doi.org/10.1063/1.1751377>.
2. H. Dreier, A. Dinklage, R. Fischer, M. Hirsch, and P. Kornejew, *Review of Scientific Instruments* **79**, 10E712 (2008), ISSN 00346748, URL <http://dx.doi.org/10.1063/1.2956962>.
3. U. von Toussaint, T. Schwarz-Selinger, M. Mayer, and S. Gori, *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* **268**, 2115–2118 (2010), ISSN 0168-583X, URL <http://dx.doi.org/10.1016/j.nimb.2010.02.062>.
4. D. V. Lindley, *The Annals of Mathematical Statistics* **27**, 986–1005 (1956), ISSN 0003-4851, URL <http://dx.doi.org/10.1214/aoms/1177728069>.
5. A. Sergeevich, A. Chandran, J. Combes, S. D. Bartlett, and H. M. Wiseman (2011), URL <http://arxiv.org/abs/1102.3700>.
6. R. Blume-Kohout, and P. Hayden (2006), URL <http://arxiv.org/abs/quant-ph/0603116>.
7. N. Boulant, T. F. Havel, M. A. Pravia, and D. G. Cory, *Physical Review A* **67**, 042322+ (2003), URL <http://dx.doi.org/10.1103/PhysRevA.67.042322>.
8. N. Boulant, J. Emerson, T. Havel, D. Cory, and S. Furuta, *The Journal of Chemical Physics* **14**, 1368+ (2004), URL <http://dx.doi.org/10.1063/1.1773161>.